

ALICE Computing Meeting
April 26, 2007

Preliminary agenda:

1. Origin of high failing rate of Alien jobs at GSI
How to proceed -
 - 1.1 Detection of problem
 - 1.2 Responsible from KP1 group to fix possible problems, if they are AliRoot related
 - 1.3 What tools can we use to detect problems(see informations below)
2. Alien SE at GSI
 - 2.1 How to make it working? - Responsible?
 - 2.2 Can we use it already for staging of the fraction of PDC06 data?
3. XRD copy
(see informations below)

=====
VERY QUICK MINUTES (27.04.2007, Silvia)

We discussed only point 1.

- a. currently only one GRID job at the time can run on the D-Grid machines, despite their 8GB memory (but only 2GB swap). Silvia will test some parallel running of AliRoot jobs executed in "stand-alone" mode instead of via AliEn (see how many jobs per machine can run, and when/where the memory limits are reached)
- b. we decided to postpone further investigations about that fact that GRID jobs cannot be executed on the batch farm to when Kilian will be back at GSI
- c. Marian and Silvia will investigate the memory consumption of AliRoot during a typical simulation+reconstruction job (using tools like valgrind). Possible problems (like memory leaks) will be reported to the software/ detector responsables.

Below, please, find quite some material collected about the memory consumption of AliRoot GRID jobs and details of the failures happening when copying to the xrootd cluster.

No formal meeting will take place in the next 2 weeks. Information and results will be distributed via the forum and emails.

=====
INFORMATION ABOUT THE GRID JOBS AT GSI

1 -- STATUS AS REPORTED BY MONALISA

1.1 Summary table for all centers: see
<http://pcalimonitor.cern.ch:8889/reports/index.jsp> (in daily, weekly, monthly summaries)

1.2 Integrated CPU time for ALICE jobs at GSI compared to Muenster, Catania and FKZ (the pledged resources of these centers are, respectively, 100, 132, 138 and 281 KSI2K)

http://pcalimonitor.cern.ch:8889/display?SiteBase=NIHAM&err=0&imgsize=800x550&interval.max=0&interval.min=2628000000&log=0&page=jobResUsageSum_time_cpu&plot_series=Catania&plot_series=FZK&plot_series=GSI&plot_series=Muenster&submit_plot=Plot&sum=0

2 -- POSSIBLE REASON 1: HIGH MEMORY CONSUMPTION

The high memory consumption by AliRoot is indicated as main reason for the low efficiency at GSI.

But:

2.1 The jobs executed at GSI do **NOT** consume a higher amount of memory compared to other ALICE GRID sites, see

Virtual:

http://pcalimonitor.cern.ch:8889/display?Nodes=max&SiteBase=Jyvaskyl&err=0&imgsize=800x550&interval.max=0&interval.min=7884000000&log=0&page=jobResUsageMax_virtualmem&plot_series=FZK&plot_series=GSI&submit_plot=Plot&sum=0

Resident:

http://pcalimonitor.cern.ch:8889/display?Nodes=max&SiteBase=Jyvaskyl&err=0&imgsize=800x550&interval.max=0&interval.min=7884000000&log=0&page=jobResUsageMax_rss&plot_series=FZK&plot_series=GSI&submit_plot=Plot&sum=0

2.2 Other ALICE GRID sites generally have machines with the same memory: at best 2GB/core (raw and swap) (information from Latchezar Betev and Jan-Fiete)

2.3 My (Silvia) experience from interactive tests (on lxial36 and lxb255) and from thousands of jobs executed on the LSF farm (batch queue) confirms the memory consumption reported by the others (this is high on an absolute scale, but not higher at GSI than other centers). The MC production jobs done on the batch farm (many with a configuration almost identical to the PDC06) over the last several months were always very successful (above 95%).

Latest numbers: Aliroot v4-05-08, PDC06 configuration file and macros, 10 events per job (max memory used during the job):

lxial36 (32 bit, 1GB mem, 2GB swap) sim: 460MB, rec 470MB

lxb255 (64 bit, 8GB mem, 2GB swap): sim 640MB, rec 600MB

It is known that the memory usage is slightly higher on 64 bit machines.

I will update the numbers for 100 events/job in short, in order to compare with the performance on the batch farm.

2 -- POSSIBLE REASON 2: GSI BANDWIDTH

[Extract from an email by Kilian, April 18, 2007]

... The I/O of the ALICE jobs (downloading input files from CERN, putting output back to CERN) rather fast outnumbers the bandwidth possibilities at GSI. This is why we are trying to

route the Grid job I/O via the "Suedhessennetz" where we have a separate 100 Mb for Grid only use for test reasons. Here I am still facing a political problem, though, that T2 packages are not yet allowed to be routed via T1 centres to the T0 centre via the dedicated T1-T0 links. But this is already known to Latchezar and the other Grid experts and we will work it out. [...]

What is not known is for WHICH FRACTION of the failures this second reason is responsible.

=====

PROBLEMS WITH COPY TO THE XROOTD CLUSTER AT GSI

Problems have been encountered and reported by Marian and Silvia when copying large amount of data to the xrootd cluster at GSI. The copy process fails at some point of the data transfer and the new copy of the file results corrupted and unrecoverable.

A few features of these failures:

- the problem affects about 50-70% of the files copied to the xrootd cluster
- it happens while WRITING to the xrootd clusters. Reading did not present problems so far
- it happens both while copying from GSI NFS disks (Silvia: MC production, both from machines on the batch farm and from /d/aliceXY disks), and from AliEn (Marian: TPC data from CERN)
- the failure is not reproducible: in repeated attempts to copy a file, the failure can re-appear or not, and usually at different times of the copy
- in Marian's example, copying TPC data from AliEn, the failure rate was about 50% when writing them to the xrootd cluster, but 0% when writing them to a local disk. See his script /u/miranov/AliRoot6/HEAD0307/TPC/macros/aliDownload.C (the 2 cases correspond, respectively, to the destinations:
TString gspath="root://lxf35.gsi.de:1094//alice/testtpc/raw2006/";
and
TString gspath="/data.local2/miranov/raw2006/");
- Marian also tried to repeat up to 5 times the copy from AliEn, but this did not improve sensibly the situation: some failure appeared all the time.

Horst suggested that the problem might be connected to poor network connections, which cause transfer delays. Extract from his email (April 24, 2007):
[...] xrootd is optimized for high performance transfers, and I guess that the time-outs in client server communication are very tough. As you possibly know our network is currently not always as stable as it should be. This could be improved with HW upgrades, but we don't have enough money for this now. I'm not an expert in xrootd internals, but if occasional network delays (~1 s) are responsible for your copy problems, repeating the copy command should help (of course, this is not acceptable, but maybe a hint to the root of the problem).[...] (about the fact that repeating the copy does NOT solve the problem, see examples above).

This should be further investigated and proved. We should see if the acceptable timeout could be set to larger times, in xrootd (if this is the reason).

One more test which is pending is copying large amount of data (TPC, PDC06) to the D-Grid machines with xrootd daemon (waiting for AliEn to come back online).

Subject: Re: Meeting on Thursday, April 26, 2007
Posted by [Kilian Schwarz](#) on Thu, 26 Apr 2007 04:39:04 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dear Silvia and all,

I would like to comment your points:

1. from the summary table one can see that GSI dedicated 100 CPUs to Grid computing, which we did not fulfill so far. But we will certainly come close to that within this year. Especially when you compute the number of 100 CPUs into kSi2k of the time when this number has been pledged.

2. you reported nothing about failure rates at GSI.

what you reported is:

- integrated CPU time at GSI compared to other centres.

This value compares more or less with the percentual site share, which has been 4 to 5% at Muenster and 0.5% at GSI so far.

The reason for this is that we have been running continuously 3 Grid jobs at GSI for quite some time (and most of them successfully and without failures !!!)

since the high mem queue did not provide more than 3 machines in the batch farm and the other batch machines were not able to deal with the memory consumption of ALICE jobs.

Since we put the D-Grid machines into production we run 10 jobs in average in parallel, which is already a significant improvement. We also tried already 2 jobs in parallel but the efficiency of job performance became definitely worse compared to only 1 job per machine. The reason might be found out by continuous local job monitoring. According to our preliminary findings the peak of memory consumption, which happens only for a short while during reconstruction, may exceed the machine capabilities when two jobs are running at the same time having the peak roughly at the same time, too.

In any case if the integrated CPU time at GSI compares with the percentual share of jobs running at GSI then the failure rate at GSI also compares with other sites.

Memory consumption of ALICE jobs at 64bit machines used to be significantly higher than at non 64bit machines with former AliRoot versions. This improved slightly with newer AliRoot versions.

If you compare the memory of GSI batch machines with other machines then please consider that our machines rarely have SWAP but mainly the on board memory. The Linux group believes that it does not make sense to use SWAP since this would slow down the jobs too much and would abuse the local disk.

The GSI bandwidth does not provide a reason for job failure. The GSI bandwidth is just limiting the amount of jobs we can run in parallel. The Suedhessennetz is not yet in production for our Grid jobs due to the mentioned political problems, which are not yet solved.

xrootd cluster: I agree with the problem analysis of Horst Goeringer and I also agree that we should follow up this more closely.

Cheers and have fun,

Kilian

Subject: Re: Meeting on Thursday, April 26, 2007

Posted by [Kilian Schwarz](#) on Thu, 26 Apr 2007 05:13:05 GMT

[View Forum Message](#) <> [Reply to Message](#)

Dear all,

because of the memory consumption of ALICE jobs:
please contact also W. Schoen. He monitored the memory consumption on the batch farm machines with some tool of his. According to my memory we had a peak consumption of something like 4-5 GB during reconstruction, I believe, and an average memory consumption of around 2 GB. For sure our finding was, that with the peak memory consumption we could not run safely two jobs at the same time on a 8 GB RAM machine, assuming that they could reach their peak within the same time slot. I discussed that many times with Walter.

Cheers,

Kilian
